



## Nor-Tech Contributes to NASA Article: Beowulf Clusters Make Supercomputing Accessible

Original article available at NASA Spinoff: [https://spinoff.nasa.gov/Spinoff2020/it\\_1.html](https://spinoff.nasa.gov/Spinoff2020/it_1.html)

### NASA Technology

In the Old English epic Beowulf, the warrior Unferth, jealous of the eponymous hero's bravery, openly doubts Beowulf's odds of slaying the monster Grendel that has tormented the Danes for 12 years, promising a "grim grappling" if he dares confront the dreaded march-stepper.

A thousand years later, many in the supercomputing world were similarly skeptical of a team of NASA engineers trying to achieve supercomputer-class processing on a cluster of standard desktop computers running a relatively untested open source operating system. "Not only did nobody care, but there were even a number of people hostile to this project," says Thomas Sterling, who led the small team at NASA's Goddard Space Flight Center in the early 1990s. "Because it was different. Because it was completely outside the scope of the supercomputing community at that time."

The technology, now known as the Beowulf cluster, would ultimately succeed beyond its inventors' imaginations.

In 1993, however, its odds may indeed have seemed long. The U.S. Government, nervous about Japan's high-performance computing effort, had already been pouring money into computer architecture research at NASA and other Federal agencies for more than a decade, and results were frustrating.

Under one Federal program, the High-Performance Computing and Communications Initiative, started in 1992, Goddard had set a goal of achieving teraflops-level supercomputing performance by 1997—a teraflops representing a speed of one trillion floating-point operations per second. And it needed to work on systems cheap enough to be assigned to a single employee. But by 1993, the best codes were achieving just a few tens of gigaflops on the biggest machine, more than an order of magnitude slower and at many times the desired cost per operation. There was no clear path to achieve performance goals within the required cost range.

Parallel processing—using a cluster of processors to achieve a single task—was understood, but each vendor was developing its own proprietary software to run on costly proprietary hardware, and none was compatible with another. A



*Thomas Sterling, who co-invented the Beowulf supercomputing cluster at Goddard Space Flight Center, poses with the Naegling cluster at California Technical Institute in 1997. Consisting of 120 Pentium Pro processors, Naegling was the first cluster to hit 10 gigaflops of sustained performance.*

single workstation might cost \$50,000 per user. The existing systems also crashed regularly, forcing reboots that restarted all running jobs and could take half an hour.

All this is documented in a 2014 paper by James Fischer, who managed the Earth and Space Science (ESS) Project at Goddard at the time. ESS housed Goddard's arm of the high-performance computing initiative, as the Center required high-speed data analysis to work with the massive amounts of data generated by its Earth-imaging satellites.

"It is only 20 years ago, but the impediments facing those who needed high-end computing are somewhat incomprehensible today if you were not there and may be best forgotten if you were," Fischer writes. In late 1992, Fischer asked Sterling if he would be interested in transitioning from NASA Headquarters to conduct supercomputing research efforts at Goddard. Around this time, Donald Becker, a longtime colleague of Sterling's, arrived in town to start a new job at another Government institution. Sterling let him stay at his place for a couple of weeks while he looked for an apartment.

"Don was very much his own thinker," Sterling says. For one thing, Becker had taken to writing code using the new, free, open source Linux operating system, which, to Sterling's thinking, was a fun toy but of little other value and sorely deficient for real work.

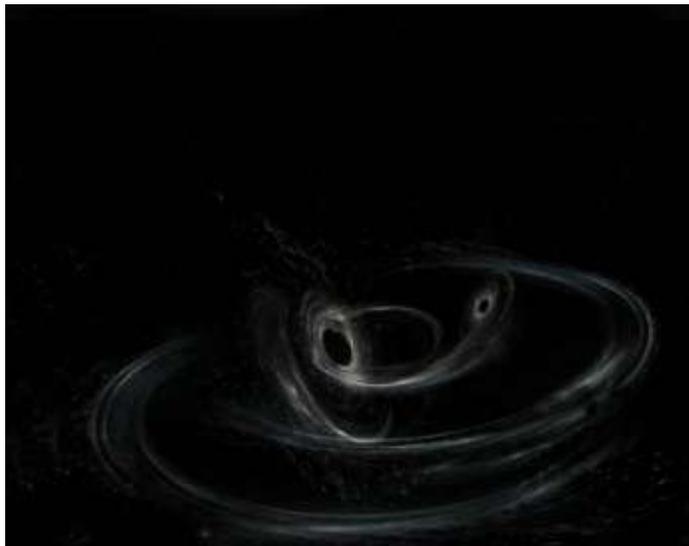
One evening Becker came home depressed and worried that he would lose his job due to a certain inattention to formalities, Sterling says. "He wasn't someone who would do the right paperwork."

This prompted Sterling's eureka moment. It occurred to him that Linux, and Becker's skill with the software, might be applied to the supercomputing problem he was grappling with at Goddard, letting him bypass supercomputer vendors entirely. "I needed a piece of software where we could get ahold of the source code, create the necessary cluster ecosystem, and provide a robust system area network at low cost," he says.

He wrote an appeal to Fischer, proposing that Goddard hire Becker to help tackle the supercomputing problem by using Linux to marshal an army of off-the-shelf personal computers to carry out massive operations. It was the first Fischer had heard of Linux, but he let himself be convinced. "He took a lot of heat, but he saw the opportunity and he didn't give in," Sterling says.

Thus NASA became the first major adopter of Linux, which is now used everywhere, from the majority of servers powering the internet to the Android operating system, to the animation software used in most major studios.

"People said Linux wouldn't last two years," Sterling recalls. "I said, I don't need more than two years. I'm just running an experiment."



*In this illustration, two black holes orbit each other, preparing to merge. Such a merger caused the first gravity waves ever detected, which the Laser Interferometer Gravitational-Wave Observatory registered in September of 2015. The Pulitzer Prize-winning project used a massive Beowulf cluster of 2,064 central processing unit cores and 86,000 graphics processing unit cores—custom-built by Nor-Tech—to crunch its data.*

By the summer of 1994, Sterling, Becker, and their team had built the first supercomputer made up of a cluster of standard PCs using Linux and new NASA network driver software. It didn't meet all the requirements of Goddard's supercomputing initiative, but subsequent generations eventually would.

Those first years were not easy, Sterling says. "It can't be done," he recalls one member of the ESS external advisory committee telling him, in remarks that were not unusual. "And if it could be done, no one would use it. And if people used it, no one would care. And if people cared, it would destroy the world of supercomputing."

That world was dominated by "people wearing suits and ties and sitting at large tables," Sterling says. "You were expected to have a large machine room and be using Fortran." The prospect of a coup by a handful of informal young people using personal computers didn't sit well. Open source freeware also was new and violated industry expectations.

"We broke all the rules, and there was tremendous resentment, and of course the vendors hated it," Sterling says. For the next 10 years, Sterling and Becker told anyone who asked that they'd named their invention for an underdog challenging a formidable opponent. But Sterling now admits this tale was invented in hindsight. In fact, he chose the name when a Goddard program administrator called and asked for a name on the spot. "I was helplessly looking around for any inspiration," he says. His mother had majored in Old English, and so he happened to have a copy of the early Anglo epic sitting in his office.

"I said, 'Oh hell, just call it Beowulf. Nobody will ever hear of it anyway.'"

### **Technology Transfer**

"If you're in the business, certainly most people have heard of it," says Dom Daninger, vice president of engineering at Minneapolis-based high-performance computing company Nor-Tech, one of a number of companies that optimize Beowulf-style clusters to meet clients' needs. The annual SC supercomputing conference even includes a Beowulf Bash every year, he says.

Despite the initial pushback, as word of the Beowulf cluster got out, interest also began to build. Sterling says even the strange name seemed to grab people's attention. In 1995, he started giving tutorials on how to build a Beowulf cluster, and two years later, he, Becker, and their colleagues received SC's Gordon Bell prize for their work. In 1999, the team co-authored the MIT Press book *How to Build a Beowulf*.

The Ethernet driver software Becker had created to overcome Linux's inherent weaknesses and tailor it to supercomputing was made open source and easily downloadable, paving the way for it to become the top operating system for supercomputing and many other applications. The creators also started the website [Beowulf.org](http://Beowulf.org), where an international community of users post questions and discuss the technology.

The difference between Beowulf and other parallel processing systems is that it has no custom hardware or software but consists of standard, off-the-shelf computers, usually with one server node controlling a set of "client nodes" connected by



*Thanks to the low-cost supercomputing made possible in part by the invention of Beowulf supercomputing clusters, manufacturers now use modeling and simulation software to design all kinds of consumer goods.*

Ethernet and functioning as a single machine on the Linux operating system. Other open source operating systems can be used. A cluster can easily be scaled up or down by adding or subtracting units.

### Benefits

"I hadn't anticipated the freedom Beowulf clusters gave a lot of people, who were no longer tied down by cost and proprietary software," Sterling says.

The technology costs anywhere from a third to a tenth of the price of a traditional supercomputer, according to Beowulf.org.

"People were free to configure their own systems, determine their software stack without interference, and contribute to the entire community," Sterling says. "Beowulf was the people's computer."

Nonetheless, it took many years to come to dominate the market.

Recently, Daninger says, many industries have come to rely on software for computer-aided engineering and modeling and simulation. But programs for such complex applications as finite element analysis and computational fluid dynamics (many of which, he notes, are based on NASA software) require huge amounts of processing and storage capacity. This used to limit their use to major corporations, such as a plane manufacturer running simulations on an airliner design before building it, saving huge amounts of money in the process.

Now the hardware for a

Beowulf cluster is cheap enough that the cost of licensing simulation software ends up being about three-fourths of the total cost of the cluster over its lifetime, Daninger says, making it affordable to smaller companies designing smaller products.

"We constantly hear from pharmaceutical and manufacturing companies that they're running out of bandwidth," he says, adding that a simulation might tie up a workstation for 72 hours before generating its first answer.

Instead, a company can use a workstation to preprocess a simulation and then offload it to a Beowulf supercomputing cluster, he says. "They can go about their work on their workstations, and then the answer will come back to them."

Nor-Tech custom-builds each cluster, integrating the servers, computers, and Linux-based operating system. Another NASA-enabled technology, the portable batch system (*Spinoff* 2001) is usually incorporated to schedule and prioritize jobs and manage resources.



*Another Pulitzer Prize-winning project, the IceCube Neutrino Observatory at the South Pole, used a Nor-Tech Beowulf cluster to process data that led to the first detection of neutrinos from beyond Earth's solar system in 2013.*

“By the time we deliver a cluster to the customer, they’ve already had a chance to remote in and test their simulation and modeling jobs, and we’ve trained them on how to use it,” Daninger says. He says customers typically start with a small cluster of four to eight units but scale up within the first year. “A lot of our customers are on their third or fourth scale-up.” Where a motorcycle manufacturer used to build 5 or 10 prototype engines and test them—an expensive and time-consuming process—now it can generate and test simulated models. “It saves them a lot of physical modeling, and there’s a quicker time to market with more reliable products,” Daninger says.

One customer, a recreational vehicle manufacturer, came to Nor-Tech after spending more than \$1 million designing, tooling, and producing a vehicle hood that it had to scrap due to heat and stress issues. A pharmaceutical company using simulation software for molecular modeling went from a three-day turnaround time to nine hours when it switched to a small cluster. “They could have gotten it down a lot more if they had scaled up,” Daninger says.

Other customers include aircraft builders running aerodynamics simulations and lawnmower manufacturers who want to predict where the machines will propel grass, among other behaviors. Companies now run modeling software on Beowulf clusters to design such everyday objects as toothpaste tubes, soap dispensers, and diapers. Others use the technology for such disparate applications as predicting beam behavior in proton-beam therapy for cancer treatment and modeling radiation emissions from a hypothetical nuclear catastrophe.

In the last few years, two of the company’s customers have won the Nobel Prize in Physics. Both the Laser Interferometer Gravitational-Wave Observatory that first detected gravitational waves and the IceCube observatory that first sensed neutrinos from beyond the solar system used Nor-Tech Beowulf clusters to crunch their data.

Nor-Tech started out as a computer parts distributor, but after almost 20 years of building clusters, these now account for about half of the 30-person company’s business.

Thanks to the Beowulf cluster, virtually every one of the world’s top 500 supercomputers runs on Linux. Four out of five of those are based on Beowulf.

“There was never a plan to revolutionize computing,” Sterling says, adding that Beowulf only triumphed because many others contributed to it over the years. “If we had tried to do something world-changing, we wouldn’t have succeeded.” Now a professor of engineering at Indiana University and president of the new start-up Simultac LLC, Sterling says, “I’m working on other things that people think are stupid—a sign that I must be doing something right.”